The route to chaos in routing games: Population increase drives period-doubling instability and chaos with Price of Anarchy equal to one

Georgios Piliouras
Singapore University of Technology and Design (SUTD)

Thiparat Chotibut	**Fryderyk Falniowski**	**Michał Misiurewicz**
SUTD	Cracow Univ. of Economics	IUPUI

Well known facts about congestion games

1. Learning dynamics such as Multiplicative Weights Updates (MWU) in congestion games converge to Nash equilibria.
2. Nash equilibria are near optimal (Price of Anarchy is small) $\Rightarrow$
   $\Rightarrow$ Via (λ, μ)-smoothness arguments regret-minimizing learning algorithms have near optimal time-average performance.

Stress tests for congestion games

1. Learning dynamics such as Multiplicative Weights Updates (MWU) in congestion games converge to Nash equilibria. In **games with many agents**, the dynamics can cycle or become chaotic.
2. Nash equilibria are near optimal (Price of Anarchy is small) $\Rightarrow$ $\Rightarrow$ Via (λ, μ)-smoothness arguments regret-minimizing learning algorithms have near optimal time-average performance. Convergence to near optimal performance can be rather slow in **games with many agents**. Practical performance can be bad even with PoA=1.

Stress tests for congestion games

1. Learning dynamics such as Multiplicative Weights Updates (MWU) in congestion games converge to Nash equilibria. In **games with many agents**, the dynamics can cycle or become chaotic.
2. Nash equilibria are near optimal (Price of Anarchy is small) $\Rightarrow$ $\Rightarrow$ Via (λ, μ)-smoothness arguments regret-minimizing learning algorithms have near optimal time-average performance. Convergence to near optimal performance can be rather slow in **games with many agents**. Practical performance can be bad even with PoA=1.

Chaos is robust and emerges even for two parallel links and linear cost functions.

Congestion games

Congestion games are a class of games in game theory first proposed by Rosenthal in 1973. In a congestion/routing game each player chooses as a strategy path in a graph (set of edges). The cost of each player depends

- on the resources he chooses
- and the number of players choosing the same resource.

Edges e.g.correspond to different driving routes.

Congestion games

A *congestion game* is defined by the tuple $(N; E; (S_i)_{i \in N}; (c_e)_{e \in E})$

- N set of players
- E set of edges/facilities/ bins
- $S_i \subset 2^E$ the set of strategies of player i
- $c_e : \{1, ..., N\} \to \mathbb{R}^+$ cost function of edge e (increasing function)

For any state $s = (s_1, ..., s_N)$

- $l_e(s)$ number of players that use edge e
- $c_i(s) = \sum_{e \in s_i} c_e(l_e)$ the cost function of player using path s

Each player aims at minimizing his cost.
There exists a (potential) function $\Phi = \sum_e \sum_{l=1}^{l_e(s)} c_e(l)$, which at each state s captures the deviation incentives for all agents:
$\Phi(s_i, s_{-i}) - \Phi(s'_i, s_{-i}) = c_i(s_i, s_{-i}) - c_i(s'_i, s_{-i}) \quad \forall s_i, s_{-i} \in S_i$
The (local) minima of Φ are Nash equilibria.

Non-atomic Congestion games

A *congestion game* is defined by the tuple $(N; E; (S_i)_{i \in N}; (c_e)_{e \in E})$

- $N \in \mathbb{R}^+$ the total system demand
- E set of edges/facilities/ bins
- $S_i \subset 2^E$ the set of strategies of player i
- $c_e : \mathbb{R}^+ \to \mathbb{R}^+$ cost function of edge e (increasing function)

For any state $s = (s_1, ..., s_N)$

- $l_e(s)$ total flow of players that use edge e
- $c_i(s) = \sum_{e \in s_i} c_e(l_e)$ the cost function of player using path s

Each player aims at minimizing his cost.
There exists a convex (potential) function $\Phi = \sum_e \int_0^{l_e(s)} c_e(l)dl$.
The (global) minima of Φ are Nash equilibria.

Price of Anarchy

Price of Anarchy (Koutsoupias, Papadimitriou '99)

The price of anarchy of a game is the ratio of social cost in the worst case scenario (assuming players are selfish) divided by the social cost of the optimum state:

$$PoA = \frac{\sup_{\text{Nash Eq.}} SC(x)}{\min SC(x)}$$

where the social cost of a state $SC(x)$ is the sum of costs of agents.

For non-atomic congestion games with linear cost functions $PoA \leq 4/3$. (Roughgarden, Tardos '00)
For atomic congestion games with linear cost functions $PoA \leq 5/2$. (Koutsoupias, Christodoulou '05)

Price of Anarchy *decreases* with many agents!

- (Feldman et al. '16) Stronger PoA bounds are possible in games with many agents. Specifically, even in the case of atomic congestion games, the better non-atomic bounds apply (i.e. $4/3$ instead of $5/2$ for linear congestion games).
- Under extra assumptions, in the heavy traffic regime PoA $\to 1$. (Colini-Baldeschi et al. '16-'17)

Simple game: Two parallel links

- agents have two available strategies $\gamma_1 = \{e_1\}$, $\gamma_2 = \{e_2\}$
- their cost functions $c_{e_1}(l) = l$, $c_{e_2}(l) = l$ depend on number (part of the flow) of players using edge/strategy e_1 and e_2 respectively.

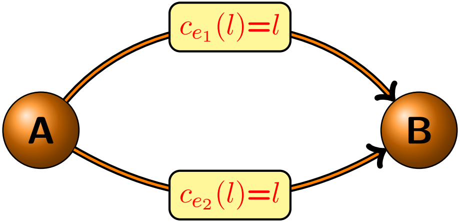

Atomic game $N = 2$, PoA= 3/2.
Atomic game $N = 2K$, PoA= 1+1/N.
Non-atomic game N, PoA=1.
Intuitively, atomic games look like non-atomic as the effects of a single agent become negligible.

Is the large population limit is enough for optimal performance?

We will study the effects a simple non-atomic congestion game with total demand N, two parallel links and linear cost functions where all agents/flow update their behavior with the same **simple learning rule**.

- PoA reduction comes at the cost of destabilizing the system. **Equilibrium is optimal**, but **learning it is trickier!**
- Every system has a carrying capacity N_0, above which dynamics is non-equilibrating.
- By increasing the total demand, we can prove that the system eventually becomes chaotic, invalidating the PoA predictions of near-optimal system performance.

Multiplicative Weights Update

A distribution for a player having the set of strategies S, is maintained on a certain set, and at each step the probability assigned to action γ is multiplied by $(1-\varepsilon)^{c(\gamma)}$, where $c(\gamma)$ is a cost of the action $\gamma \in S$, $\varepsilon \in (0,1)$ is learning rate

$$p_\gamma(t+1) = p_\gamma(t) \frac{(1-\varepsilon)^{c_\gamma(t)}}{\sum_{\gamma' \in S} p_{\gamma'}(t)(1-\varepsilon)^{c_{\gamma'}(t)}}$$

MWU is the discrete time variant of replicator dynamics.
It is known to have good properties from the perspective of online optimization.

Regret

We compare the expected reward of an algorithm to the reward incurred by the *best fixed action* in hindsight.

Definition

Fix cost vectors $\mathbf{c}_1, \ldots, \mathbf{c}_T$. The (expected) regret of the (randomized) algorithm $\mathcal{A}$ choosing actions according $x_1, \ldots, x_T$ is

$$\underbrace{\sum_{t=1}^{T} E_{s_n \sim x_n} c_n(s_n)}_{\text{our algorithm}} - \underbrace{\min_{s \in S} \sum_{n=1}^{T} c_n(s)}_{\text{best fixed action}} . \tag{1}$$

The regret of MWU with parameter ϵ given any sequence of costs $\vec{c}_1, \ldots, \vec{c}_T \in [0,1]^m$ is at most $\epsilon * T + \ln(m)/\epsilon$.
In the case of two actions, $m = |S| = 2$, the regret of MWU is at most $\epsilon * T + \ln(2)/\epsilon$.

Full model: Simple congestion game + MWU(ϵ)

Non-atomic congestion game with two edges e_1, e_2 and total demand N.

- All agents update their probability distribution $(x, 1-x)$ using **MWU with fixed learning rate** ϵ. Total flow on e_1 is Nx, whereas on e_2 it is $N(1-x)$.
- Cost functions are linear functions $c_{e_1}(l) = \alpha l$, $c_{e_2}(l) = \beta l$ where l is part of the flow (players) using edge/strategy e_1 and e_2 respectively.

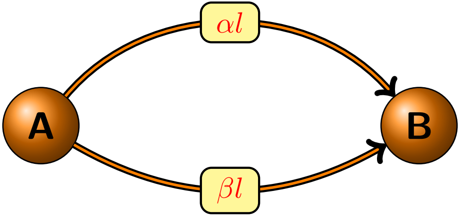

One dimensional dynamical system due to MWU applied on a simple congestion game

Dynamics which MWU generates can be described by $f_{a,b}\colon [0,1] \mapsto [0,1]$, where

$$f_{a,b}(x) = \frac{x}{x + (1-x)\exp(a(x-b))},$$

with $a > 0, b \in (0,1)$, and

$$a = (\alpha + \beta) N \ln\left(\frac{1}{1-\varepsilon}\right),\ b = \frac{\beta}{\alpha + \beta}.$$

Dynamics introduced by the MWU algorithm

Dynamics which MWU generates can be described by $f_{a,b}\colon [0,1] \mapsto [0,1]$, where

$$f_{a,b}(x) = \frac{x}{x + (1-x)\exp(a(x-b))},$$

with $a > 0, b \in (0,1)$, and

$$a = N \ln\left(\frac{1}{1-\varepsilon}\right),\ b = \beta$$

with $\alpha + \beta = 1$.

Transformation $f_{a,b}(x) = \frac{x}{x+(1-x)\exp(a(x-b))}$

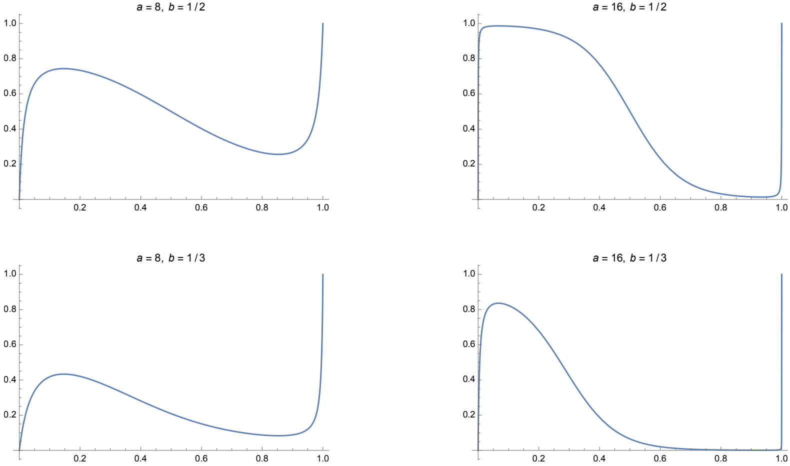

Fully symmetric case — same cost functions

Theorem

Let $c_{e_1} = c_{e_2}$, that is $b = \frac{1}{2}$.

1. *If $0 < a \leq 8$ then $f_{a,1/2}$-trajectories of all points of $(0, 1)$ converge to the fixed point $1/2$ (Nash equilibrium).*
2. *If $a > 8$ then $f_{a,1/2}$ has a periodic attracting orbit $\{\rho_a, 1 - \rho_a\}$, where $0 < \rho_a < 1/2$. This orbit attracts trajectories of all points of $(0, 1)$, except countably many points, whose trajectories eventually fall to the repelling fixed point $1/2$.*

Fully symmetric case — same cost functions

Theorem

Let $c_{e_1} = c_{e_2}$, that is $b = \frac{1}{2}$.

1. *If $0 < a \leq 8$ then $f_{a,1/2}$-trajectories of all points of $(0,1)$ converge to the fixed point $1/2$ (Nash equilibrium).*
2. *If $a > 8$ then $f_{a,1/2}$ has a periodic attracting orbit $\{\rho_a, 1-\rho_a\}$, where $0 < \rho_a < 1/2$. This orbit attracts trajectories of all points of $(0,1)$, except countably many points, whose trajectories eventually fall to the repelling fixed point $1/2$.*

Proof hint: By Sharkovsky's theorem any discrete dynamical system $f : [0,1] \to [0,1]$ where f is a continuous function if it does not have a periodic point of period 4 it only has fixed points and periodic points of period 2. By symmetries of $f_{a,1/2}$ we can argue that every initial condition either converges to a fixed point or a periodic orbit of period 2. When $a > 8$ the Nash equilibrium becomes unstable.

Broken symmetry — cost functions are different

What about asymmetric case, when cost functions are different???

CHAOS

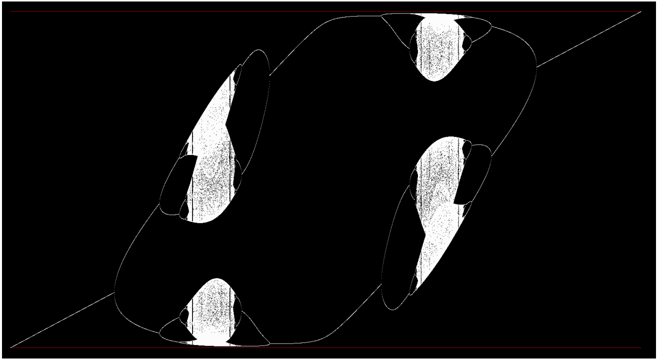

The route to chaos

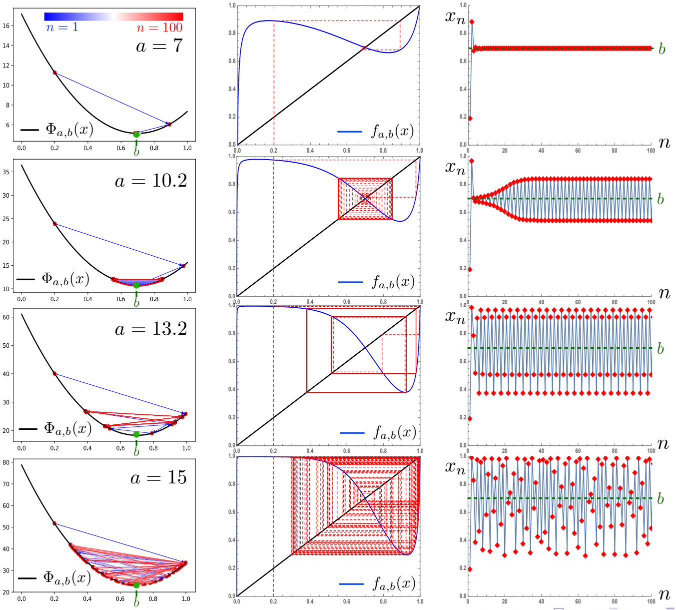

Chaos for the interval map

Li-Yorke chaos

Let (X, f) be a dynamical system and $(x, y) \in X \times X$. We say that (x, y) is a *Li-Yorke pair* if

$$\liminf_{n\to\infty} dist(f^n(x), f^n(y)) = 0,$$
$$\limsup_{n\to\infty} dist(f^n(x), f^n(y)) > 0.$$

A dynamical system (X, f) is Li-Yorke chaotic if
there is an uncountable set $S \subset X$ (called scrambled set) such that every pair (x, y) with $x, y \in S$ and $x \neq y$ is a Li-Yorke pair.

Topological entropy for maps

For each natural number n, a new metric d_n is defined on $X = [0,1]$ by the formula

$$d_n(x,y) = \max\{d(f^i(x), f^i(y)) : 0 \leq i \leq n\}$$

Given any $\delta >$0 and $n \geq 1$, two points of X are δ-close with respect to this metric if their first n iterates are δ-close.
A subset E of X is said to be (n,δ)-separated if each pair of distinct points of E is at least δ apart in the metric d_n.
Denote by $N(n,\delta)$ the maximum cardinality of an (n,δ)-separated set. The topological entropy of the map f is defined by

$$h(f) = \lim_{\delta \to 0} \left(\limsup_{n \to \infty} \frac{1}{n} \log N(m,\delta) \right)$$

Different cost functions always lead to chaos for large enough rates

Theorem

Let $\alpha \neq \beta$ *(i.e.,* $b \neq \frac{1}{2}$*), where* b *isthe Nash equilibrium:*

- *For* $a \in (0, 4)$ *Every trajectory converges to Nash equilibrium.*
- *If* $b \in (0, 1)\backslash\{1/2\}$*, then there exists* a_b *such that if* $a > a_b$*, then* $f_{a,b}$ *has periodic orbits of all periods, positive topological entropy and is Li-Yorke chaotic.*

The route to chaos

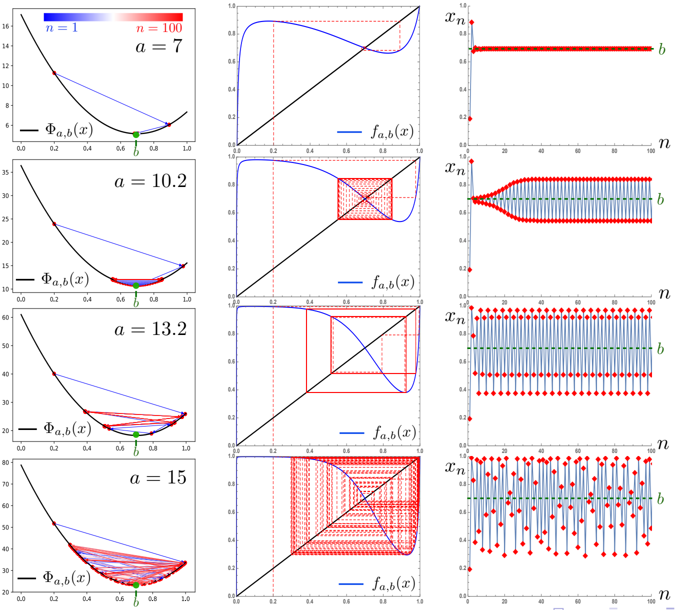

Average behavior — like in zero-sum games

For an interval map f a point p is Césaro attracting if it has a neighborhood U such that for every $x \in U$ the averages

$$\frac{1}{n}\sum_{k=0}^{n-1} f^k(x)$$

converge to p.

Theorem

For every $a > 0$, $b \in (0,1)$ and $x \in (0,1)$ the Nash equilibrium b is Césaro attracting for $f_{a,b}$.

Corollaries

Corollary

For every periodic orbit $\{x_0, x_1, \ldots, x_{n-1}\}$ *of* $f_{a,b}$ *in* $(0,1)$ *its center of mass (time average)*

$$\frac{x_0 + x_1 + \cdots + x_{n-1}}{n}$$

is equal to b*.*

Corollary

For every probability measure μ*, invariant for* $f_{a,b}$ *and such that* $\mu(\{0,1\}) = 0$*, we have*

$$\int_{[0,1]} x \, d\mu = b.$$

Game theoretic implications

We know that

$$a = N \ln\left(\frac{1}{1-\varepsilon}\right), \quad b = \beta$$

To simplify calculations we put $\varepsilon = 1 - \frac{1}{e}$, then $a = N$.

a – the (normalized) system demand
b – (normalized) equilibrium flow.

Corollaries

- If an interior equilibrium is not $50\% - 50\%$ split, increasing the total demand of the system will inevitably lead to chaotic behavior, regardless of the form of the cost function.
- Time-average flows on the edges converge exactly to equilibrium values — a property akin to no-regret learning in zero-sum games.

And what about regret and social costs?

Regret and variance

Regret in our case:

$$R_T = \sum_{t=1}^{T} (\alpha N x_n^2 + \beta N (1 - x_n)^2) - \min\{\sum_{n=1}^{T} \alpha N x_n, \sum_{n=1}^{T} \beta N (1 - x_n)\}$$

Theorem

The time-average regret is the total demand N times the variance of the variable $x_n = f^n(x)$

$$\lim_{T\to\infty} \frac{R_T}{T} = N \left(\lim_{T\to\infty} \frac{1}{T} \sum_{n=1}^{T} (x_n - b)^2 \right) = N \cdot \mathrm{Var}(X).$$

Time-average social cost and variance

$$\mathrm{Var}(X) = \lim_{T\to\infty} \frac{1}{T} \sum_{n=1}^{T} (x_n - b)^2$$

We can relate the normalized time-average social cost to the variance. We have

$$\begin{aligned}
\text{norm. time-average soc. cost} &= \frac{\text{time-average social cost}}{\text{optimum social cost}} \\
&= \frac{\frac{1}{T}\sum_{n=1}^{T}\left(\alpha N^2 x_n^2 + \beta N^2 (1-x_n)^2\right)}{N^2 \alpha\beta} \\
&= \frac{\frac{1}{T}\sum_{n=1}^{T}\left(x_n^2 - 2\beta x_n + \beta\right)}{\beta(1-\beta)} \\
&= 1 + \frac{\mathrm{Var}(X)}{\beta(1-\beta)}.
\end{aligned}$$

Upper bound for variance

Remark

The variance is bounded above by

$$\mathrm{Var}(X) = \lim_{T\to\infty}\left(\frac{1}{T}\sum_{n=1}^{T}(x_n - b)^2\right) \leq (f_{\max} - b)(b - f_{\min}).$$

where $f_{\max}, f_{\min}$ *the largest and smallest respectively value of the map at its critical points.*

Regret, social cost and chaos

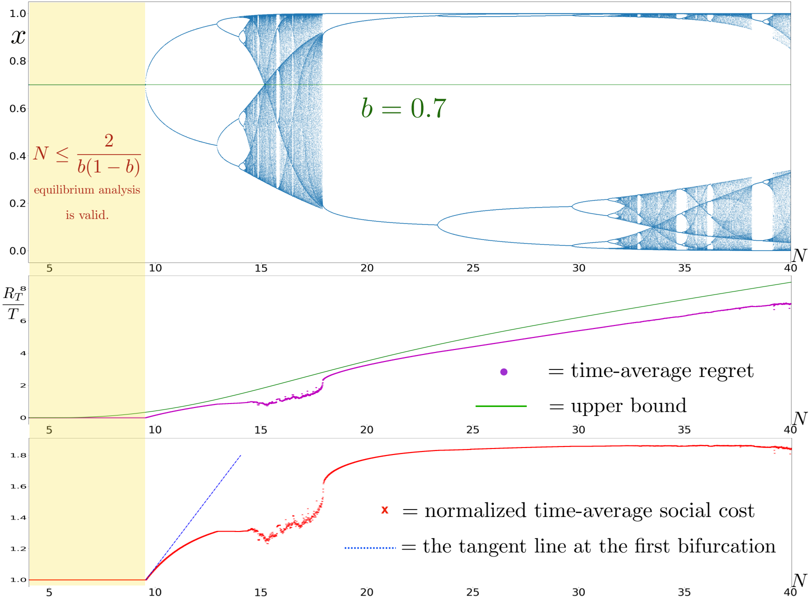

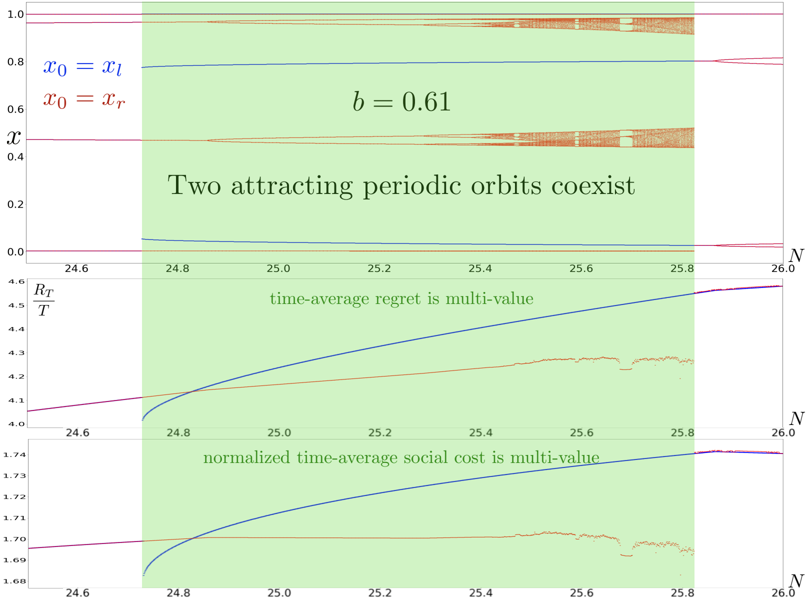

$x_0 = x_l$
$x_0 = x_r$
$b = 0.61$
Two attracting periodic orbits coexist
x
N
$\frac{R_T}{T}$
time-average regret is multi-value
normalized time-average social cost is multi-value

Extensions and problems

Extensions

- More strategies
- Heterogeneous case with two different learning rates
- Different dynamics

Extensions and problems

Extensions

- More strategies
- Heterogeneous case with two different learning rates
- Different dynamics

Problems

- Provable chaos in other games (e.g. zero-sum games)?
- Variance of the interval map (even for the logistic map)
- Lower bound for regret (variance), sensitive dependence on parameters?

Upper bound

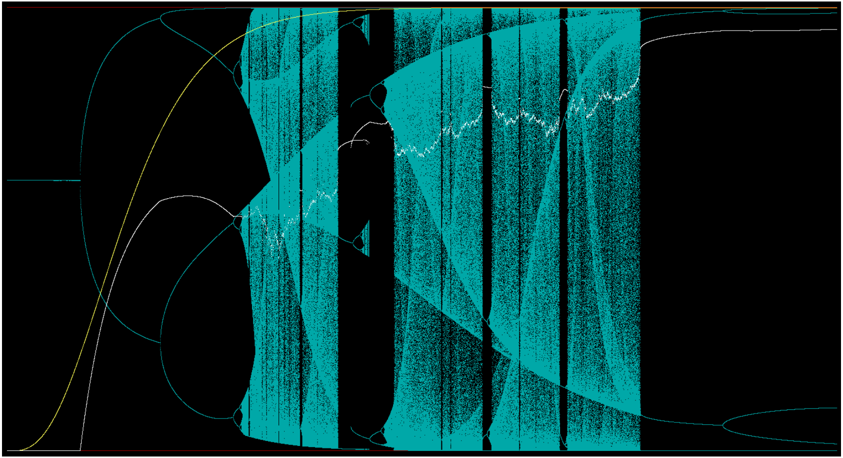

Figure: Attractor (cyan) and average regret (white) with the bound (yellow) for $b = 0.61$, $a \in (4, 54)$

Upper bound

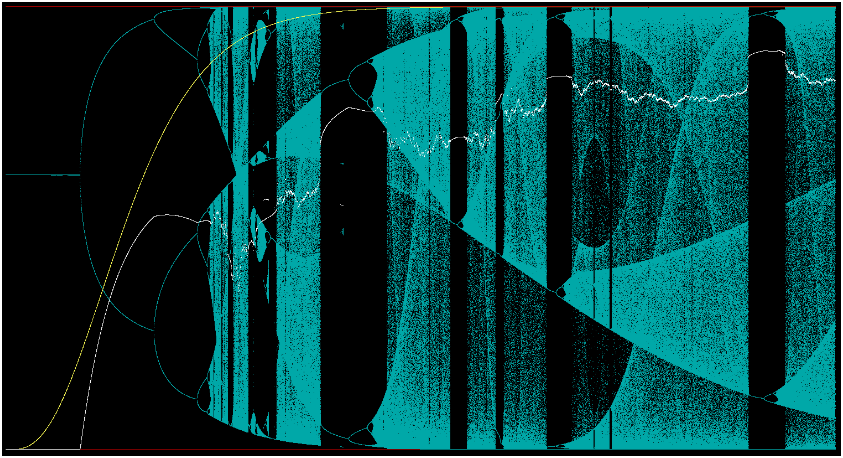

Figure: Attractor (cyan) and average regret (white) with the bound (yellow) for $b = 0.62$, $a \in (4, 54)$

Upper bound

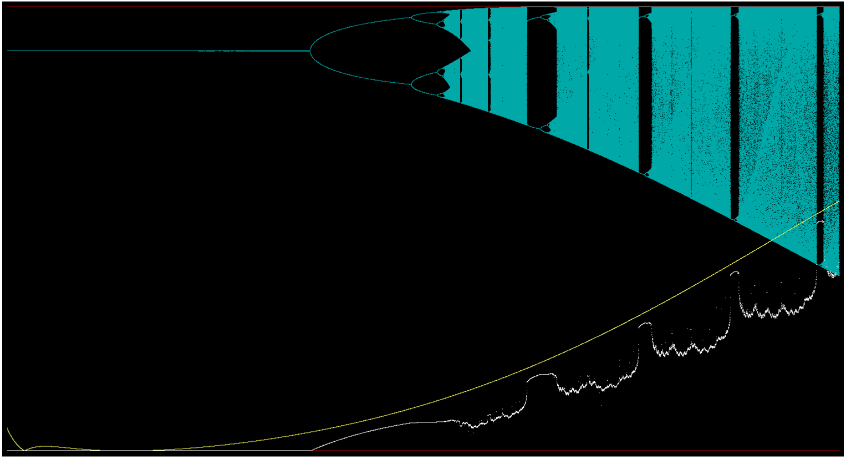

Figure: Attractor (cyan) and average regret (white) with the bound (yellow) for $b = 0.90$, $a \in (4, 54)$

Proper upper bound

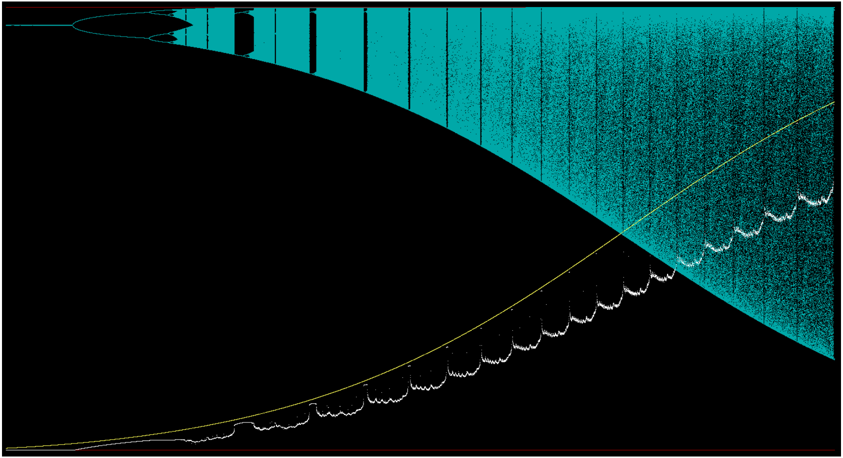

Figure: Attractor (cyan) and average regret (white) with the bound (yellow) for $b = 0.96$, $a \in (40, 190)$

Proper upper bound

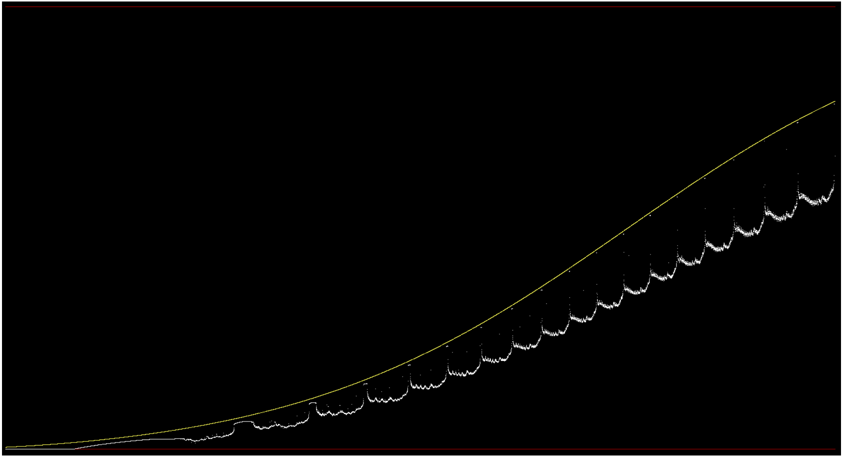

Figure: Average regret for $b = 0.96$, $a \in (40, 190)$

Proper upper bound: $b = 0.96$ zoom

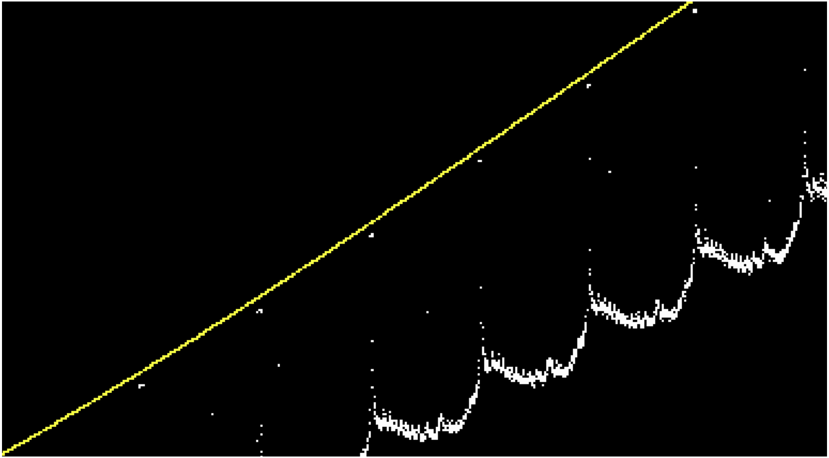

Thank you!